



Europäisches Patentamt
European Patent Office
Office européen des brevets



⑪ Publication number:

0 631 229 A2

⑫

EUROPEAN PATENT APPLICATION

⑬ Application number: 94108769.4

⑮ Int. Cl. 5: G06F 9/44

⑭ Date of filing: 08.06.94

⑯ Priority: 14.06.93 US 77348

⑰ Date of publication of application:
28.12.94 Bulletin 94/52

⑱ Designated Contracting States:
DE FR GB

⑲ Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504 (US)

⑳ Inventor: Erickson, Charles Robert
1509 Spiderlily View
Cedar Park,
Texas 78613 (US)
Inventor: Sessions, Roger Hereward
11414 Pencewood Drive
Austin,
Texas 78750 (US)

㉑ Representative: Lettieri, Fabrizio
IBM SEMEA S.p.A.,
Direzione Brevetti,
MI SEG 024,
P.O. Box 137
I-20090 Segrate (Milano) (IT)

㉒ Method and system of storing and restoring object persistent attribute data.

㉓ A system and method for storing object information on a persistent medium. A DefaultEncoderDecoder is supplied that encodes object persistent attribute data into a string of attribute names, types and data values. The encoding and decoding is managed by one or more get and one or more set methods. The get methods control storing of the data while set methods control restoring of persistent data. Encoded strings of attribute data are upwardly compatible through replacement or overriding of the set methods.

EP 0 631 229 A2

The present invention relates to data processing systems and more particularly to systems for creating application programs or operating system programs. Still more particularly, the present invention relates to object oriented systems having persistent object storage functions operable to store and restore persistent objects without operator modification.

5 The development of application and system software for data processing systems has traditionally been a time consuming task. The field of software engineering has attempted to overcome the limitations of traditional techniques by proposing new, more efficient software development models. Object oriented programming has emerged as a promising technology that will allow rapid development, implementation and customization of new software systems.

10 Object oriented programming uses a toolkit of system objects that can be assembled to perform the final task. Each object has certain data attributes and processes or methods that operate on that data. Data is said to be "encapsulated" by an object and can only be modified by the object methods. Methods are invoked by sending a message to an object identifying the method and supplying any needed arguments.

15 Object oriented systems have two important properties in addition to encapsulation. "Inheritance" is the ability to derive a new object from an existing object and inherit all properties, including methods and data structure, from the existing object. The new object may have certain unique features which are supplied as overrides or modifications to the existing class. I.e. a new subclass needs to specify only the functions and data that distinguish that class from the existing more general class.

20 The ability to override an existing method description enables polymorphism, the ability to have a single message to an object be processed in different ways depending on the object itself.

Inheritance and polymorphism create a powerful structure for implementing new software systems. The software developer does not have to develop each piece of a system, he or she need only specify the unique features of the system.

25 The power of object oriented systems is realized through the development of system "frameworks." A framework is a collection of base classes that can be used by a system implementor to create a final systems product. The framework is defined and developed to work together as a system. Conceptually, the framework is much like a set of standard hardware components used by computer hardware builders. Each of the components has certain defined functions and interfaces and the engineer assembles these components according to a particular design to create a unique hardware system.

30 Objects created by a data processing system are typically maintained in the volatile memory of those systems. This allows faster processing, but does not provide a means to store data. Object oriented systems solve this problem by implementing "persistent objects." "Persistence" means that an object's state can be preserved beyond the termination of the process that created that object. The persistent framework include methods that allow persistent objects to be stored in and retrieved from a non-volatile medium such as a hard disk drive or writable optical disk. Objects can be stored individually or grouped with other objects.

35 The use of Persistent Streams to capture and retain database queries and results is discussed in EP-A 90480133.9, filed on 4/9/90 and commonly assigned (U.S. Patent No. 5,161,225).

40 Existing persistent object systems are limited by their fixed implementation of object storage. A particular object oriented product typically provides one mechanism for storing persistent objects. This mechanism includes a definition of the non-volatile medium used, the structure of object groups stored on the medium and the structure of each particular object stored on the medium.

45 Fixed implementations of persistent objects limit the developer flexibility to change persistent object storage to meet particular system requirements. A fixed storage scheme requires that the method for restoring data remains constant thereby constraining the developers from improving or enhancing the application. A change to the number or types of parameters to which an object responds would prevent restoration of previously stored objects in a fixed storage system.

50 Fixed implementations have the advantage, however, of requiring no coding effort by the developer before being able to store persistent objects. The developer frequently does not have the time or training to be able to develop a customized persistent storage framework. These developers need an ability to store persistent data in a flexible framework without the concerns of later system incompatibility and without the need to devise their own storage mechanism.

55 The existing technical problem addressed by the present invention is to provide a persistent object framework that provides flexible and upwardly compatible persistent object storage without requiring developer coding of the storage methods.

The above drawbacks of the prior art are overcome by the invention as claimed.

It is therefore an object of the present invention to provide a persistent object system that enables persistent storage and retrieval of objects without requiring programming by the operator.

It is yet another object of the invention to provide an object framework that allows restoration of persistent objects without storing persistent object format information with the object.

It is still another object of the invention to provide persistent data storage that is upwardly compatible using object oriented techniques.

5 The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawing wherein like reference numbers represent like parts of the invention, wherein:

Figure 1 is a block diagram of system on which the present invention operates;

Figure 2 is an illustration of an object according to the present invention;

10 Figure 3 is a hierarchical representation of classes in the preferred embodiment;

Figure 4 is a representation of a persistent data storage group skeletal framework;

Figure 5 is a flowchart depicting the DefaultEncoder process according to the present invention;

Figure 6 is a flowchart depicting the DefaultDecode process according to the present invention.

The preferred embodiment will be described with reference to the above figures.

15 An object oriented programming system operates on a computer system such as that shown in Figure

1. The computer system can be a "Personal Computer", workstation, minicomputer or mainframe type of computer. Figure 1 illustrates a personal computer such as the IBM PS/2 Model 95 computer (IBM and PS/2 are trademarks of the IBM Corporation.) The computer typically includes a system unit 100 which contains a processor or CPU 102 and memory 104. An object oriented program 106 operates on the

20 processor and is resident in memory 104. The system will typically have a display device 108 such as a color or monochrome CRT, a keyboard 110, a pointing device 112 such as a mouse or trackball, a fixed disk drive 114, and optionally, a network connection 116 to a LAN or other network.

An object encapsulates data and the methods needed to operate on that data. Objects can be represented by a "doughnut diagram" such as shown in Figure 2. Object data is shown in the center 202 surrounded by the applicable methods 204 to 214. Data 202 may be modified only by the methods of that object. Methods 204-214 are invoked by receiving messages from other objects. A typical object oriented system will have a message router 220 that routes messages between objects. Thus, object 230 causes Method 208 to be invoked by sending a message 232 to the message router that in turn sends it 222 to Method 208 of object 200.

30 Object frameworks are constructed to provide a set of objects for application and system developers to use to construct a delivered system. The IBM System Object Model (SOM) framework, for example, provides a language independent set of objects for use in systems development. The IBM OS/2 Version 2.0 Workplace shell is an example of an application developed using such an environment.

Objects are grouped into classes of related objects. The class description contains information relevant to all objects in a class, including a description of instance variables maintained by each of the objects and the available object methods. An object instance is created (or "instantiated") based on that information and has the properties defined in the object class. For example, the object class DOG can include the instance variables "dog_type" and "dog_name" and a "bark" method implementing the response to a bark message. An instance of dog, e.g. ROVER, will maintain the type and name instance variables for itself and will respond to the bark message.

40 Abstract classes are used to describe the interfaces and methods expected to be used by a class without providing detail on the implementation of those methods. Abstract classes are useful in frameworks where the implementation details are to be left to the implementor. Concrete classes are created as subclasses of abstract classes and implement those classes.

45 Classes in an object oriented system are often derived from some single base class. This base class is named "SOMObject" in the IBM System Object Model. Figure 3 illustrates the preferred embodiment of the present invention by showing the class hierarchy for Persistent Objects using the System Object Model.

The class "PersistentObject" 304 is a subclass of SOMObject 302 and contains the methods and variables necessary to create persistent objects. Because it is a subclass of SOMObject it inherits all 50 methods and variables known to SOMObject. Also derived from SOMObject 302 are PersistentStorageManager 306, IOGroupManager 308, and EncoderDecoder 310. These subclasses are abstract classes for implementing overall management of persistent objects (PersistentStorageManager), for controlling the structure and grouping of persistent objects on the persistent object storage medium (IOGroupManager), and for writing object information to the storage medium (EncoderDecoder).

55 The IOGroupManager abstract class 308 controls the grouping and organization of groups of objects on the persistent storage medium. Persistent objects are stored in related groups of objects. The structure of object group storage is called the Group Skeletal Format and defines a skeleton within which the objects are stored. The most straightforward structure is storing the object description in plain ASCII text with

objects stored sequentially within a group. Figure 4 illustrates the group skeletal format for this object structure.

The group skeletal format 400 contains a Object Group Header 402 that contains group data, a number of object separators 404 406 ... 412 and object storage blocks 414 416 ... 420. Storage of objects within a storage block 414 is controlled by the EncoderDecoder class.

EncoderDecoder 310 functions to reformat object instance data from the internal object storage format to a format that can be written to a persistent storage medium (e.g. a disk.) The persistent format must contain sufficient information to restore the object to internal storage in the same state as when it was written.

An EncoderDecoderAbstract class is implemented in the preferred embodiment to define the interfaces to concrete EncoderDecoder subclasses derived by the application developer or developer. The EncoderDecoderAbstract class defines the interface to the following methods. EncoderDecoderWrite formats and writes the instance data of a particular object using a selected media interface. EncoderDecoderRead reads from the selected media interface and restores the object from the persistent format to internal object format.

The EncoderDecoder of the present invention sends messages to the selected IOGroupManger for all media specific functions, including preparing the device, opening and closing the device, and physically writing and reading the data from the device.

The preferred embodiment of the present invention provides a DefaultEncoderDecoder to implement flexible, generalized storage of object information. The DefaultEncoderDecoder relies on a repository of information about the class of an object including the data types of attributes of the class. It further relies on a means to identify attributes of a class as persistent attributes. The repository used to contain the above information about an object is known as an Interface Repository. Finally, the DefaultEncoderDecoder relies on a "set" and a "get" method for each persistent attribute, a consistent mapping of persistent attribute name to the name of the method used to get/set the attribute and a consistent get/set method signature for each supported data type. The "get" method or methods control storage of the persistent object, while the "set" method or methods control restoration of the object.

The preferred embodiment implements attributes as specified by the standards of the Object Management Group. Attributes are the public data associated with objects. Each attribute is required by the standards to have a get and set method. Persistent attributes are indicated by the keyword "persistent" next to the attribute names in the Interface Definition Language (IDL) entries for the object.

The DefaultEncoderDecoder is activated when storage of an object is required. The DefaultEncoderDecoder consults the Interface Repository to determine if the object has any data attributes and if so whether any of the data attributes have been identified as persistent. For each of the persistent object data attributes, the DefaultEncoderDecoder consults the Interface Repository to determine the data type of the attribute. Given the name of the persistent attribute and its data type, the DefaultEncoderDecoder constructs a "get" method request. The attribute "get" method is then invoked to retrieve the data to be stored. The DefaultEncoderDecoder generates a stream of data for each persistent data attribute. The stream of data includes the data type of the persistent attribute, the attribute name and the data returned from the attribute "get" method. All basic attribute data types are supported including integers, floating point numbers, character strings, arrays and sequences of these types and structures of any of these types. The resulting stream is concatenated with all streams generated from that object and written to persistent storage.

The persistent stream, according to the present invention, is written as an ASCII stream of attribute names and parameters. Although the preferred embodiment writes an ASCII stream, any other format could be used, including EBCDIC or binary. The stream is in a readily parsed format to allow reconstruction of the object.

The structure of the object can be understood by reference to the following example:

```

interface Person : SOMPPersistentObject
{
5      attribute string name;
      attribute short deptNum;

10     // Attribute get/set methods:
15     void _set_name(string name);
           string _get_name();

15     void _set_deptNum(short num);
           short _get_deptNum();

20     implementation
{
25       // Attribute modifiers identify which
           // attributes are persistent
           name: persistent;
           deptNum: persistent;
30   };
35 };

```

The DefaultEncoderDecoder stores the essential object information as an ASCII string having the
35 following preferred format:

1. For each persistent attribute in the object:
 - a. number of characters in the persistent attribute name;
 - b. name of attribute;
 - c. data type of the attribute;
 - d. data returned by the attribute "get" method.
 - 40 2. An "End of Attributes" marker to signal the end of the data stream upon restore of the object.
- The stream of data stored for an instance of the above example object would result in a string such as:

45

50

55

```
(4)name(19)David(7)deptNum(20)(987)EOA
```

```

      |   |   |   |   |   |   |   |   |
      |   |   |   |   |   |   |   |   |   End of attrib marker
5    |   |   |   |   |   |   |   |   |   data for attrib "deptNum"
      |   |   |   |   |   |   |   |   |   data type for attrib "deptNum"
      |   |   |   |   |   |   |   |   |   name of attribute
10   |   |   |   |   |   |   |   |   |   length of next name of attribute
      |   |   |   |   |   |   |   |   |   data for attribute "name"
      |   |   |   |   |   |   |   |   |   data type for attribute "name"
15   |   |   |   |   |   |   |   |   |   name of attribute
      |   |   |   |   |   |   |   |   |   length of next name of attribute

```

The process of the preferred embodiment for encoding object data for persistent storage will be described with reference to Figure 5. The Persistent Object Manager 306 is instructed by the program to store an object. The PersistentStorageManager in turn instructs an IOGroupManger 308 to instantiate an EncoderDecoder, the DafaultEncoderDecoder according to the present invention, and the DefaultEncoderDecoder invokes the necessary get methods.

The process starts at 502 and proceeds to reading the Interface Repository to determine the number and types of data attributes for the class of this object 504. A test for the existence of more attributes is made at 506. If no more attributes exist an "end of attributes" marker is stored 508 and the process terminated 520. If more attributes exist, a test is made to determine if the next attribute is persistent 510. If not, processing returns to test 506. If a persistent attribute is found, the length of the attribute name is stored 512 and the attribute name and type are stored 514. A "get" method request is constructed using the attribute name and data type and the method invoked 516. Finally, the data returned from the "get" method is stored 518 and processing returns to test 506.

Objects are restored from persistent storage through the invocation of set methods for the persistent attributes of the object. The request to restore an object is initiated by the program by invoking a method on the Persistent Storage Manager. The Persistent Storage Manager then instantiates an IOGroupManager which in turn instantiates an EncoderDecoder (in this case the DefaultEncoderDecoder.)

The restoration of an object is accomplished by the process steps illustrated in Figure 6. The process starts at 602 and proceeds immediately to step 604 that reads an attribute name from the persistent storage stream. The result is tested to determine whether or not it is an "end of attribute" marker 606. If so, the restoration ends 616. Otherwise, the system reads the attribute type 608, allocates space to hold the restored data 610, and reads the data stream into the allocated space 612. A "set" method is constructed using the attribute name and data type and the method is then invoked 614. Upon return, control passes to step 604.

This storage format allows the implementation and even the interface to the object to be changed without loss of persistent data. Upward compatibility of data is controlled by overriding the "set" methods to support the current structure.

Claims

1. A system for storing data for a selected object to a persistent medium in an object oriented computer system also having a memory, a processor, and operator interaction means, wherein data resides in system objects in said memory, and wherein said data is manipulated by processing methods associated with said object, the system comprising:
 - attribute extraction means for determining a number, names, and types of attributes defined for said selected object, said attribute extraction means being responsive to a data storage request;
 - method construction means for constructing a method invocation request to get attribute values for the attributes extracted by said attribute extraction means, said method construction being responsive to said data storage request;
 - data formatting means for formatting said attribute names, types and values for storage to said

- 10 persistent medium; and
15 data restoration means for restoring said attributes to a restored object instance.
- 20 2. The system of claim 1, wherein said method construction means invokes a "get" method.
- 25 3. The system of claim 1 or 2, wherein said data restoration means comprises:
30 persistent medium access means for reading data from said persistent medium;
35 storage allocation means for allocating storage in said memory for attribute data to be restored;
40 method construction means for constructing a method invocation request to restore attribute data.
- 45 4. The system of claim 3, wherein the method construction means constructs a "set" method request
50 using attribute name and type.
- 55 5. A method of storing and restoring object persistent attribute data to and from a persistent storage
60 medium, said method operating in an object oriented system wherein each object instance contains
65 attribute data that is manipulated by process methods, and wherein information on the number and
70 type of attributes of an object is stored in a repository, the method comprising the steps of:
75 defining a plurality of data extraction methods for each persistent attribute;
80 accessing said repository to determine the number, name and type of persistent attributes for said
85 object;
90 constructing a request to invoke one of said plurality of data extraction methods based on said
95 name and type of persistent attributes to return a plurality of attribute data values;
100 storing to said persistent medium a string of attribute names and types;
105 storing to said persistent medium a string of said attribute data values;
110 restoring said object data by successively executing restoration methods associated with said
115 attributes.
- 120 6. The method of claim 5 wherein said restoring step includes the steps of:
125 reading an attribute name from said persistent medium;
130 determining attribute type from said persistent medium;
135 allocating space for said attribute value;
140 constructing a method invocation request to set the restored attribute data.

35



40

45

50

55








FIG. 3




FIG. 4




FIG. 5

